September-October 1995
Volume 2, Number 12

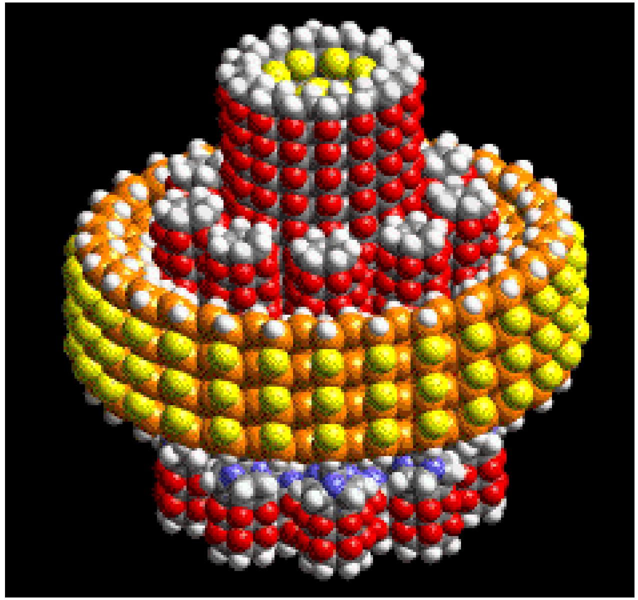

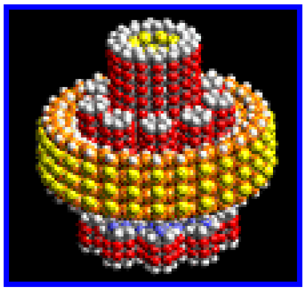

Nanotechnology Creates New Opportunities for NAS, Ames, Industry

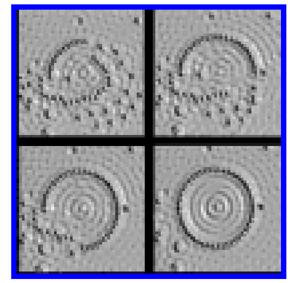

by Elisabeth Wechsler

Molecular nanotechnology -- the building of products by placing atoms in precise locations using controlled trajectories -- is "a promising future technology," said David Bailey, NAS senior scientist, in his introduction to a New Technology seminar on this subject at the NAS Facility, on July 28.

Scientists Al Globus and Creon Levit, of the NAS information management group, told the audience of some 50 NASA Ames researchers that the use of nanotechnology could achieve "orders-of-magnitude improvements in aerospace vehicles and computers." The speakers alternated presentations of what they envision nanotechnology can offer the aerospace industry -- and the role NAS could play.

Bailey, Globus, and Levit acknowledged the "substantial risk and long lead time (5-20 years)" for this field, but they nevertheless believe nanotechnology "appears feasible" and recommended that NAS "aggressively seek a dominant role in computational nanotechnology."

Computational Nanotechnology

One area particularly well-suited to NAS is the "crucial" area of computational nanotechnology, in which molecular editing, simulation, modeling, testing, and validation are used to design "molecular products with millions of atoms and thousands of assembly steps," Levit said.

According to Bailey and Globus, potential applications of computational nanotechnology include:

- ultrastrong components for jet engines
- ultralight materials for space vehicles
- computer processors with dramatically higher performance and lower energy requirements
- computer memory devices with dramatically higher storage densities

Levit explained that physical experiments for nanotechnology are often impossible and that simulation is necessary to establish design feasibility. Product complexity requires molecular Computer Aided Design (CAD) and databases of parts, properties, and processes.

He noted that NAS is well-positioned to become "the key computational nanotechnology center due to our computer systems and expertise in parallel algorithms, simulation, visualization, and operations." In addition, Levit pointed out Ames' expertise in computational chemistry and materials science, which can contribute to the inherently multidisciplinary nature of nanotechnology.

Proximity to Stanford University, IBM Almaden Research Center, Xerox Palo Alto Research Center, Foresight Institute, and Silicon Valley provide an excellent local environment, Levit continued, adding that NAS-based work would complement that done by private industry.

Fits Ames' Mission Statement

In keeping with the Ames mission statement (designating Ames as the NASA Center of Excellence for Information Systems Technologies), Globus believes that NAS could, at a minimum:

- provide the "enormous computer power necessary."
- fund university research through NASA Research Announcement grants.
- develop some of the CAD/CAM software for molecular nanotechnology.

The next step in the definition of this new area will be to sponsor a computational nanotechnology workshop.

For more information about nanotechnology work at NAS, email **globus@nas.nasa.gov** or **creon@nas.nasa.gov**.

Next Article | Contents | Main Menu | NAS Home

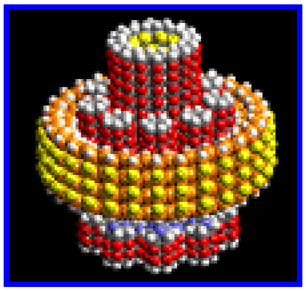

Molecular model of a hypothetical planetary gear, used in mechanical engineering, in which the spheres represent individiual atoms. This simulation suggests that moderately complex mechanical parts can be built from precisely positioned atoms. The image was designed and simulated by K. Eric Drexler (Foresight Institute, Palo Alto, CA) and Ralph Merkle (Xerox Palo Alto Research Center). It was rendered using Molecular Assembly Sequence Software (MASS), developed by Carol B. Shaw.

to the article.

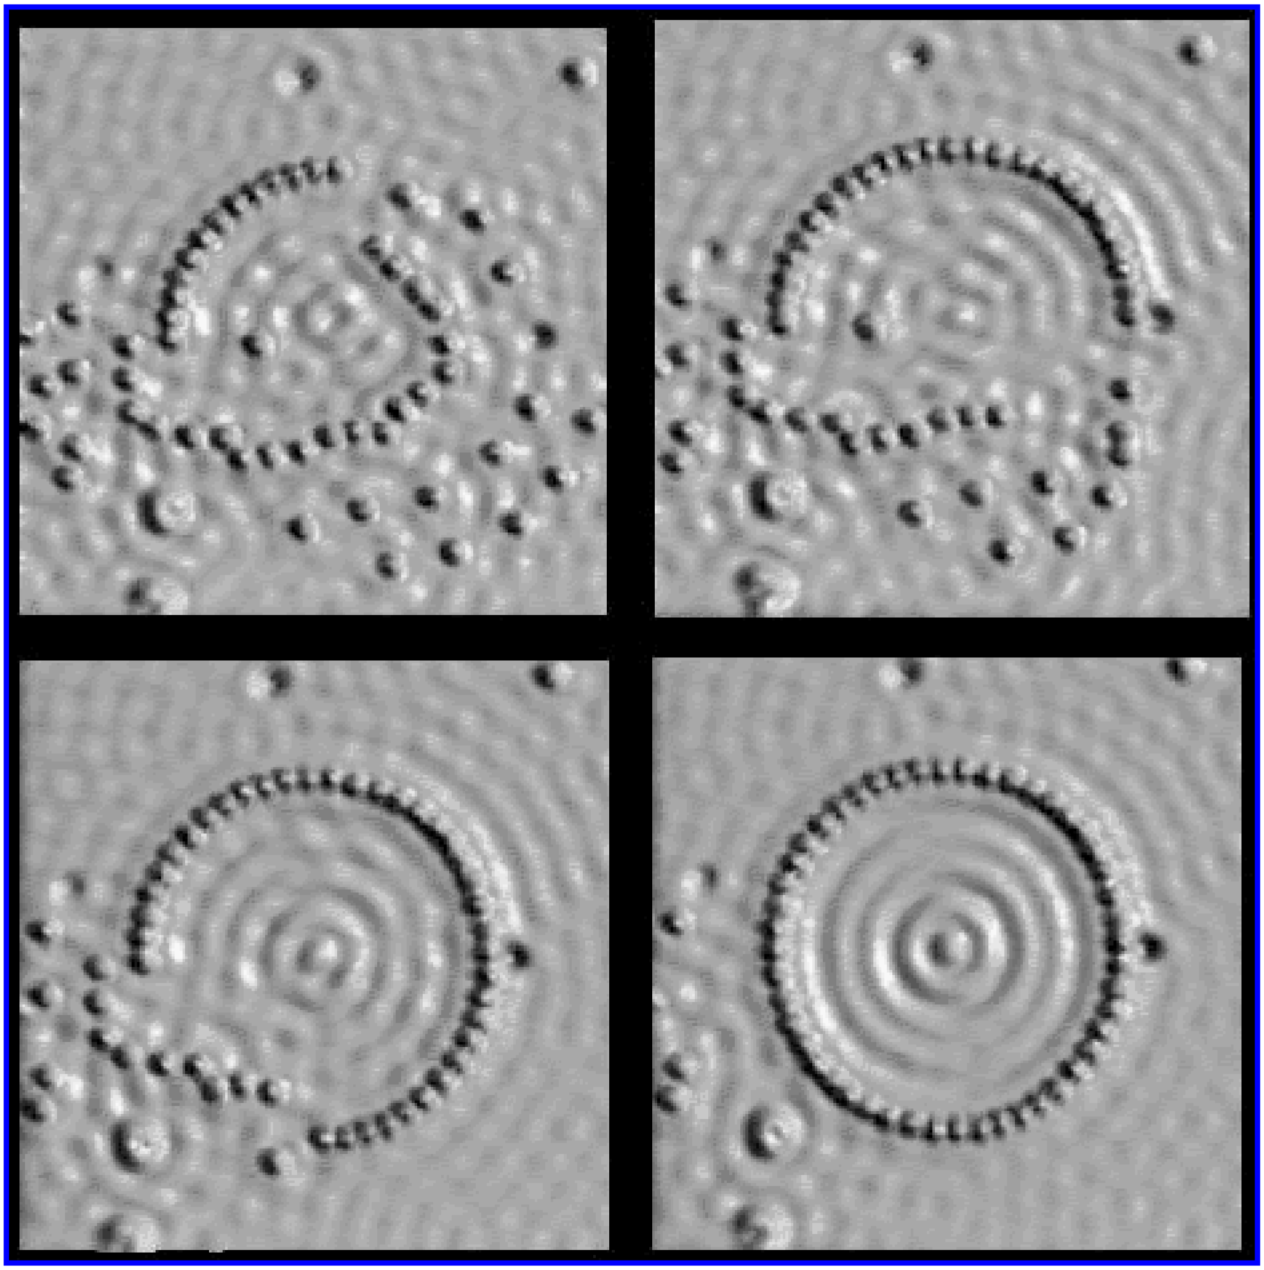

Individual atoms being "herded" into a "quantum corral." This experiment represents "one small step towards building extremely strong materials (75 times stronger than steel) and extremely small computers (10^19 MIPS on a desktop) with atomic precision," explained Al Globus, NAS scientist. The image was created using a scanning tunneling microscope by M. F. Crommie, C. P. Lutz, and D. M. Eigler at IBM Almaden (CA) Research Center.

to the article.

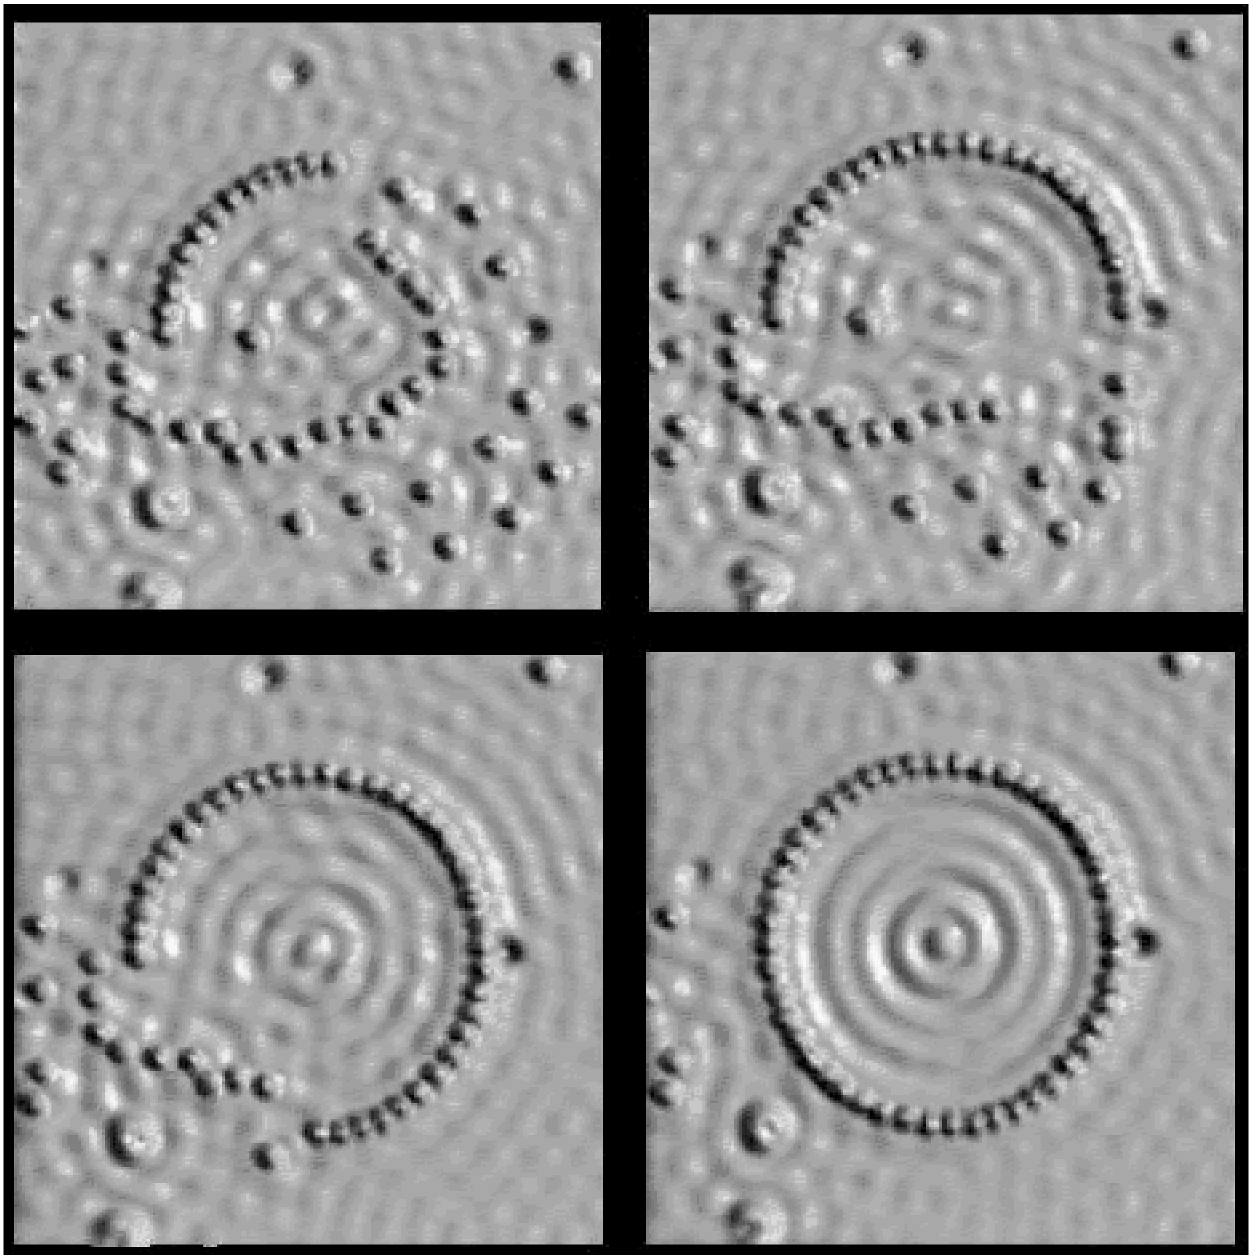

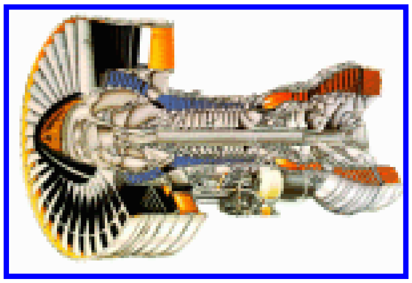

Large-scale Model for Analyzing Aircraft Engine Design Challenges Supercomputers at NAS Facility

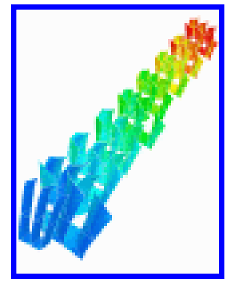

by Elisabeth Wechsler

The supercomputers housed at NAS -- eagle [CRAY C90, 8 CPUs], vonneumann [CRAY C90, 16 CPUs], and babbage [IBM SP2, 160 nodes] -- are being used this summer to run a quasi-production code that attempts to analyze the core compressor of an aircraft engine.

The mathematical model, referred to as the Average-Passage flow model, describes the three-dimensional flow in multistage turbomachines, and was developed by John Adamczyk, senior aerospace scientist, and his team at NASA Lewis Research Center.

This work has had "a significant impact on the PW4000 High-Pressure Compressor development effort" at Pratt & Whitney, a United Technologies Corp. subsidiary, according to Mark Barnett, senior research engineer at United Technologies Research Center. The model allows engineers "to obtain much more accurate predictions of the performance and characteristics of their multistage turbomachinery designs," he added.

So far, the NAS Facility's SP2 has been used to run a simulation for 17 (out of a possible 23) blade rows of a Pratt & Whitney high-pressure compressor (HPC). Adamczyk's goal is to run the full 23 blade rows in parallel "by the end of the summer."

Reduces Design Time, Cost

The use of Adamczyk's model within Pratt & Whitney's HPC analysis code has contributed to a 2 percent performance improvement over the baseline design with increased stall margins, and helped to reduce the design time from two years to one.

The potential cost savings could be "in the tens of millions of dollars" through reductions in the number of compressor rigs that would need to be built and tested, Barnett commented.

"Engine manufacturers are waiting in line to use this code," said Jim Craw, Acting Chief of the NAS Scientific Computing Branch. Others using Adamczyk's model include General Electric Aircraft Engine, Evandale, OH; Allison Engine Co., Indianapolis, IN; and Allied Signal Engines, Phoenix, AZ.

While Barnett cautions that there are still challenges "to properly include the effects of unsteadiness in multistage turbomachinery," he applauds Adamczyk's continuing effort, which provides "the potential for further significant impact on Pratt & Whitney's ability to design turbomachinery in less time and at lower cost."

Maximizing Parallelism's Benefits

For the NAS technical support staff, the biggest challenge for a long-running job such as Adamczyk's model is "a stable system with a large amount of memory," commented Chuck Niggley, group lead for NAS scientific consulting.

The Average-Passage model's unique inherent parallelism was exploited by Adamczyk's group by running the code in the multitasking queue on the high-speed processors, but -- as the project's allocations on eagle and vonneumann became exhausted -- the SP2 was harnessed. The scientific consulting group helped the SP2 system administrators set up a special file system "to guarantee continuity and efficient end results."

"This project has challenged us to think about giving time to other forms of parallelism besides multitasking on vonneumann," said George Myers, also of the NAS scientific consulting group. "We're struggling with: What's the best way to use the system? What makes sense?"

At press time, Adamczyk told NAS News that a simulation of the entire 23 blade rows on the SP2 had been completed. Results are now being analyzed by Lewis researchers, he said.

Next Article | Contents | Main Menu | NAS Home

Large-scale Model for Analyzing Aircraft Engine Design Challenges Supercomputers at NAS Facility

Cutaway of a high bypass ration turbofan engine with pointer to high pressure compressor.

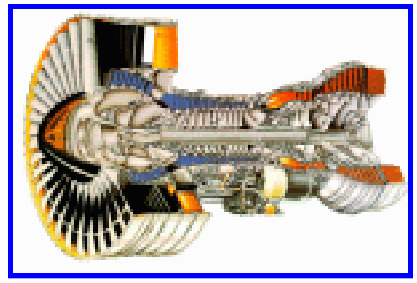

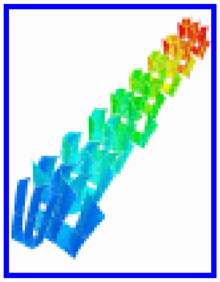

Simulation of the Pratt & Whitney 4000 compressor atr design operating condutions using the Average Passage flow model developed by NASA Lewis researchers. Pressure distributions on each blade surface with the compressor are shown.

to the article.

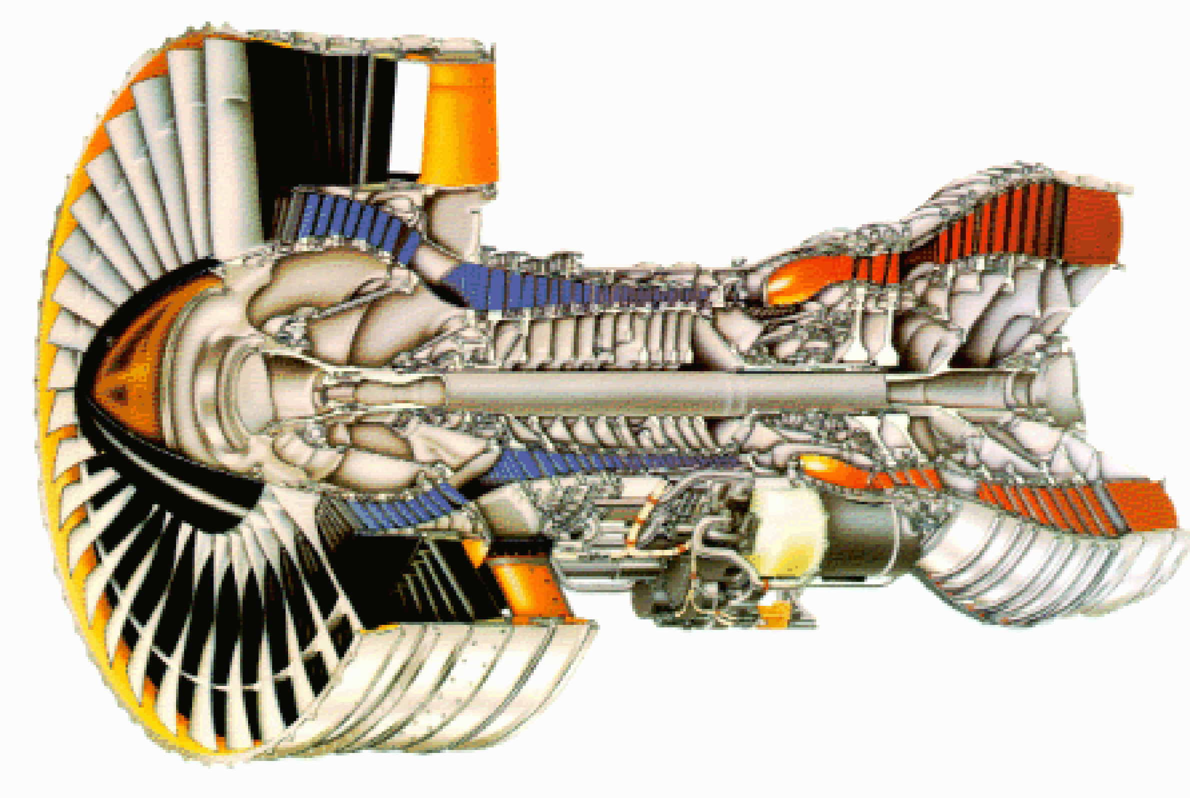

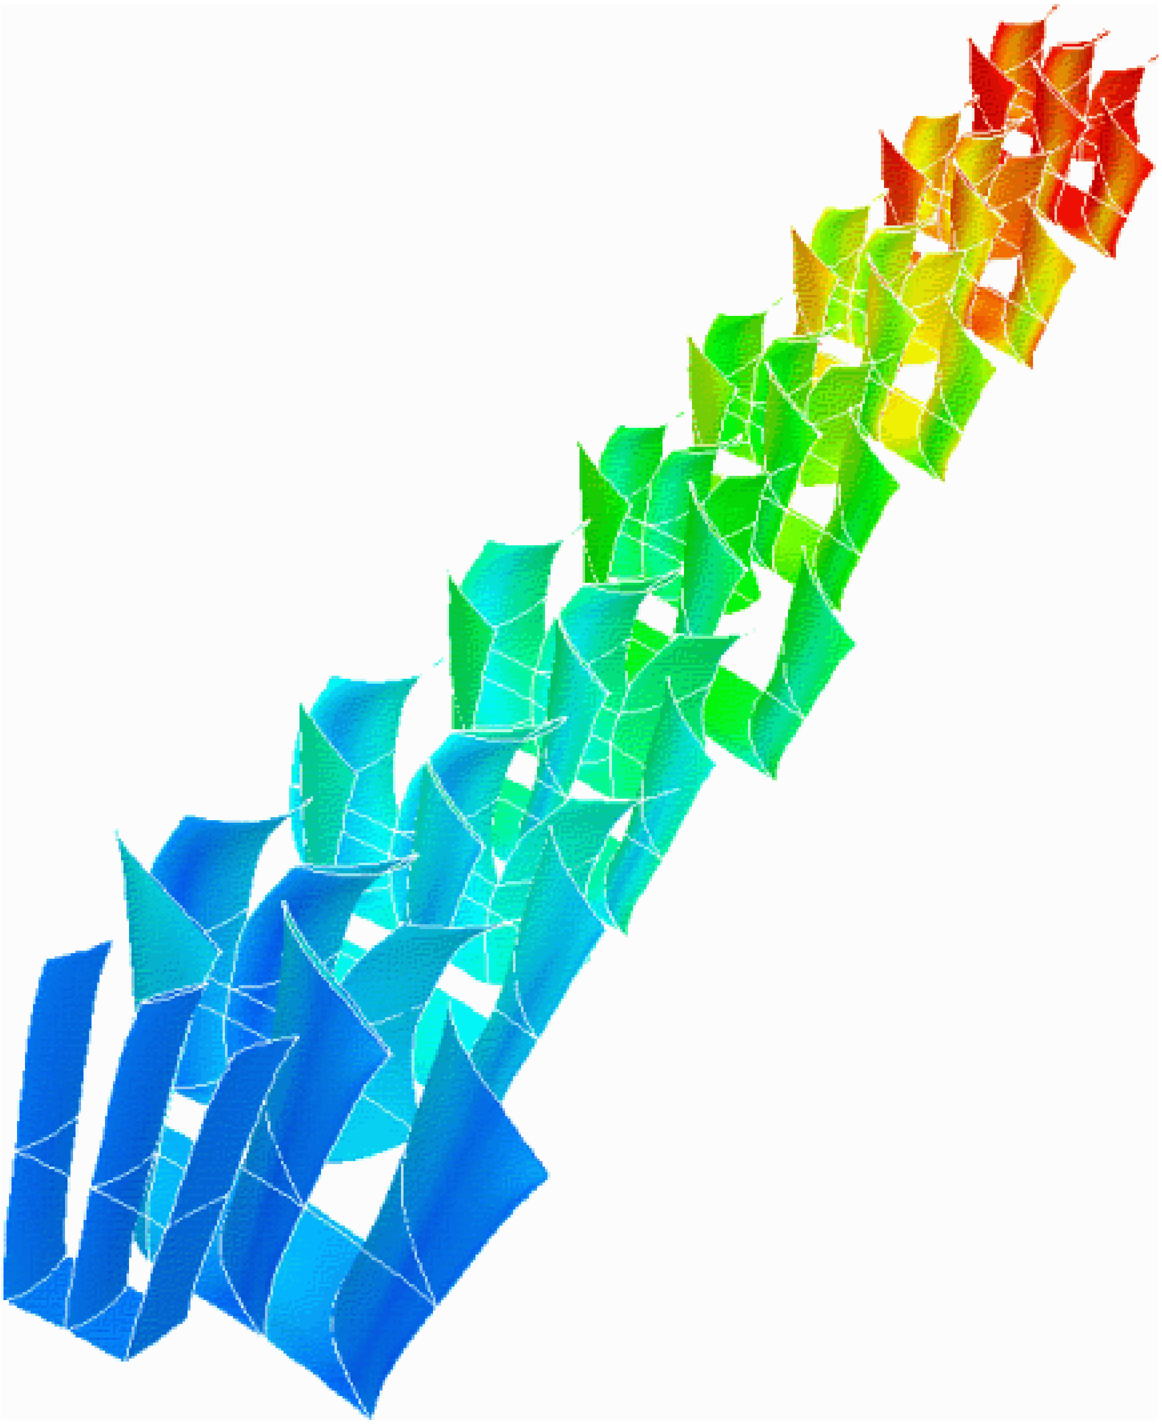

Remodeled NAS `Vis Lab' Offers Users the Latest in Interactive Data Analysis and Presentation Capabilities

by Elisabeth Wechsler

The NAS Visualization Lab is nearing completion of its nine-month remodeling project. Both local and remote users, the latter of which comprise about 40 percent of the "Vis Lab's" total customer base, will be able to do more interactive data analysis with new and upgraded tools and equipment. In addition, data presentation capabilities have been substantially expanded, and a staff of expert consultants is available to assist NAS users with their visualization needs.

Vis Lab group lead Kevin McCabe is very enthusiastic about the new facility. "A major design goal was to make the lab's configuration flexible -- to allow it to grow and change as visualization requirements change." The remodeling project, which cost approximately $500,000, is expected to be completed this month.

More Space, Improved Infrastructure

Major physical improvements to the Vis Lab include a redesigned floor plan with dedicated work areas; improved networking and data storage infrastructure; remote-controlled track lighting; a production version of the NAS Virtual WindTunnel (VWT); an upgraded video production system; and a new installation of the Stereo Visualization Theater, which was demonstrated in the NAS booth at Supercomputing '94.

The interactive data analysis tools include the VWT, developed by Steve Bryson, computer scientist; acoustics analysis software written by Dennis DeRyke, also of the NAS Vis Lab (with input from McCabe, a systems engineer, and Akil Rangwalla, an aeronautical engineer at Ames); and data analysis software such as FAST [Flow Analysis Software Toolkit].

The room housing the Vis Lab (inside the NAS building) measures 48-ft. by 26-ft. An 80-in. by 60-in.

screen with two Electrohome Marquee 9000 projectors (which comprises the Stereo Visualization Theater) highlights a presentation area near the entry with seating capacity of up to 30. This area can be used by researchers to demonstrate CFD work to their peers and serve as a site for small visualization-based workshops, as well as for acquainting visitors with the NAS Facility.

Serving Users' Requirements

On a recent tour of the Vis Lab, McCabe explained that output from any workstation in the lab can be displayed on the rear projection screen for large-scale viewing. VHS, U-Matic, and Betacam tapes can be played with audio from any source in the theater, and high-quality videotape recordings can be made simultaneously during live demonstrations.

Across the room from the theater is the new production version of the NAS VWT. Output from the VWT can be switched to the Stereo Visualization Theater, which will display in stereo what the researcher is manipulating using the boom interface. Also available are two high-end Silicon Graphics (SGI) workstations used for scientific visualization and a non-linear video editing system for multimedia projects.

Enclosed areas of the Vis Lab serve as a recording studio for computer animations; a small sound-proof booth for recording voice-over narrations; and a suite for combining animations, audio, and video.

An equipment room houses the lab workstations, air conditioning unit, HiPPI and ethernet connections to mass storage, and video equipment, including a high bandwidth routing system to reduce the need for rewiring as well as to improve the lab's infrastructure.

Providing Additional Capabilities

Some examples of what the Vis Lab will be able to do include:

- traditional video editing with multiple tape decks, as well as computer-based non-linear editing
- digital multi-track audio capability that adds voice-over narration to visual results (such as for animations), with software to accompany and support interactive analysis of acoustics in CFD datasets
- co-visualization (that is, researchers interactively collaborating with remote colleagues) using FAST simultaneously at each site
- static and time-dependent data can be analyzed by generating animations, which are recorded to digital video disks or traditional video tape machines, enabling playback of simulations in real time

- a stereophonic acoustics library, which can be linked to visual simulation programs to make acoustic data audible, and a stand-alone prototype in which users wear headphones to gain spatial perception of acoustic phenomena in simulation data

- a 600 dpi Sharp JX 610 scanner, which can capture digital images from hard copy, transparencies, and 35mm slides

Sharing Visualization Knowledge

"We want to invite researchers to contribute ideas on multimedia and visualization techniques that they've developed, and work together with our customers to resolve new problems," McCabe said.

The Vis Lab staff, available to provide consultation on visualization problems for NAS users, also includes Chris Gong, video production and publication graphics; Vee Hirsch, video production and system integration; and DeRyke, acoustics software engineering and video production support with FAST.

For more information about the NAS Visualization Lab, send email to **vis-lab@nas.nasa.gov**.

Next Article | Contents | Main Menu | NAS Home

The newly installed Stereo Visualization Theater offers a presentation area with seating for up to 30 in the remodeled NAS Visualization Lab. This site can be used to demonstrate CFD work, such as this ascent transonic flow field dataset by Fred W. Martin Jr. and other researchers at Johnson Space Center and Lockheed Engineering and Sciences Co.

to the article.

NAS Support Staff Helps Users 24 Hours a Day

by Jill Dunbar

Just two weeks from the new operational period for the Aeronautics Consolidated Supercomputing Facility (ACSF) on October 1, the NAS control room staff is gearing up for the usual onslaught of calls from new users whose problems range from "Why can't I log in?" to "Can you help me set up my environment to make the best use of my allocation?"

The control room is the central point for reporting, tracking, and resolving problems on NAS systems -- and, according to Chris Kleiber, NAS control room manager, 90 percent of all problems are resolved or have a workaround in one day. The 14 computer analysts who answer questions and research problems for both new and experienced users are available 24 hours a day, seven days a week to support users on the East Coast and many local researchers who work late or early hours.

The control room is particularly involved with new users three times a year, during the new operational periods for the parallel system, which began August 14, the NAS CRAY C90 (vonneumann), which begins in March, and the ACSF CRAY C90 (eagle). The staff works with user interface manager Pat Elson to ensure that the entire application process, which involves mounds of paperwork, database entry, and user notification, runs smoothly.

What To Expect When You Call

When a user calls the control room, a "real" person answers the phone -- no recordings, no touch-tone choices. "We attempt to answer any question," Kleiber said. They succeed 60-75 percent of the time -- usually immediately.

What happens with the calls that the control room can't answer? Depending on the nature of the problem (usually those concerning hardware, operating systems, or FORTRAN), these calls are transferred to a scientific consultant specializing in high-speed processors (HSP) or parallel processors, or to the NAS workstation or network groups through the NAS Call Management System, which immediately sends email to the consultants. For these problems, users get a response within 24 hours.

Some areas that the control room staff can address directly include:

- local workstation environment problems, incorrect permissions, full file systems

- Cray utilization, including selecting the correct file system for special uses
- NAStore utilization, efficient file storage and retrieval
- network connectivity problems for remote users
- quota, disk, and CPU problems
- shell, UNIX, and C language questions

The staff works with users in many other areas, including: enhancing NQS [Network Queuing System] scripts; answering IBM SP2; configuration questions; helping with electronic mail and editor problems; restoring files on all platforms; and providing FrameMaker support. The group also has expertise in visualization tools, including PLOT3D and NAStran.

Users need to be aware of the differences between local and off-site support, Kleiber emphasized. Control room support is available to users of NAS and ACSF systems: vonneumman, the SP2, NAStore (the NAS mass storage file system), and support processing systems and workstations located at the NAS Facility.

Off-site ACSF users, including "local" users at Ames, are often confused about the types of support the NAS control room offers, Kleiber said. Those users should call their facilities' help desks for anything other than CRAY C90 (eagle) issues, including general UNIX, and local mass storage and workstation support.

Biggest Benefit is Behind the Scenes

The control room's biggest benefit to users, in Kleiber's opinion, is the "behind-the-scenes stuff." The control room staff "has its eyes on the system all the time," she said. "If a job is hanging in the middle of the night, we can push it through so it doesn't stagnate until the next morning."

One example of this type of support occurred last spring, during a 25-hour dedicated job on the SP2 for Charbel Farhat, a researcher at the University of Colorado, Boulder (*see " Large, Out-of-core Calculation Runs on The IBM SP2*," NAS News, *July-August '95*).

Mary Hultquist, lead systems analyst for the SP2, relied heavily on control room support during this special run. She requested many last-minute complicated changes, and the control room staff responded quickly and reliably, creating new scripts, coordinating with all those involved in the project, and bringing the system back up when Farhat's job finished. Although there were "minor glitches along the way," Hultquist said that "the entire endeavor was an overwhelming success."

"What's apparent is that this group works hard to make life easier for users," Hultquist added.

"We're kind of the hub for all information and details of activities" for the NAS computer environment, Kleiber said. During a recent unexpected power outage that left all of NASA Ames Research Center without electricity -- including network connections to remote user sites -- the control room coordinated communications among building engineers, hardware field engineers (FEs), and NAS management and technical staff to track progress and report status. Working alongside FEs and systems analysts, the control room staff also helped return all systems to production within four hours after power was restored.

Metrics Bring Improvement

What's apparent from talking with Kleiber is that this group works hard to make life easier for users -- and that they're always looking for ways to improve their skills and knowledge. Kleiber tracks the group's performance against NAS standards -- and when performance falls below the "acceptable" range, she is accountable to NAS management.

For example, at one point when the number of user questions and problems resolved by the group declined, Kleiber investigated and found that the group wasn't prepared for a system upgrade. As a result, the staff now tests all upgrades before they're released to users, and stays involved with other NAS technical groups' work -- particularly the HSP and parallel groups.

"Assisting users in getting their jobs through the NPSN [NAS Processing System Network] with a minimum of difficulty is our number one priority," she said.

Get Toll-Free Help

For assistance with NAS systems, call the toll-free hot line at **1-800-331-8737** (USER) or locally call **415-604-4444**. You can also request help online.

Next Article | Contents | Main Menu | NAS Home

Taking Advantage of New Features on the Parallel Cluster

By James P. Jones

In mid-July, the High Performance Computing and Communication (HPCC) Program workstation cluster at NAS received an upgrade that replaced part of the original cluster configuration (*for details, see* NAS News, *May-June '95*) with new hardware. Along with the upgrade, changes were made to the cluster's usage model, allowing users to make more choices about -- and interact with -- their jobs.

davinci : The Latest Renaissance

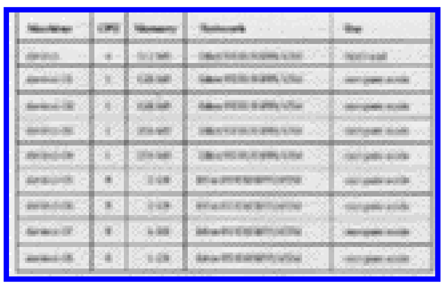

The new workstation cluster consists of one front-end system and eight "compute nodes." The front end (named "davinci") is a Silicon Graphics Inc. (SGI) PowerChallenge L with four 75-megahertz (MHz) R8000 CPUs and 512 megabytes (MB) of memory. This is where users can log in and create, submit, or debug their parallel applications. It is also the cluster file server and PBS [Portable Batch System] server. The configuration for the eight compute nodes is shown in the table.

All machines are connected via Ethernet, FDDI, HiPPI, and ATM [Asynchronous Transfer Mode].

The davinci cluster's usage model is similar to that of the NAS IBM SP2 in three ways: users perform interactive work from the front-end system, not the compute nodes; users submit jobs via the PBS queuing system; and each node will be dedicated to one user at a time -- that is, the system will be space-shared rather than time-shared among users. Though not normally used on workstation clusters, this concept of dedicating nodes to one user for the duration of a job was implemented after proving successful on the SP2.

One major difference between the workstation cluster and the SP2 is that IBM's Parallel Operating Environment ("poe"), which facilitates running parallel jobs across multiple nodes, is not available on the cluster. So, users must take an extra step to run parallel applications on the cluster using commands such as mpirun, pvmexec, or rsh in the place of poe. For detailed information on using these techniques, see "How To Use MPI on the Cluster" or " How To Use PVM on the Cluster" on the NAS parallel cluster home page.

The NAS parallel systems groups are currently investigating ways to add functionality similar to poe to the cluster. (Send suggestions via email to **davstaff@nas.nasa.gov**.)

As with the NAS SP2, PBS is the job queuing system for the cluster. Users submit jobs via PBS, specifying the resources needed, such as the number and type of nodes; the amount of wall-clock time needed; and memory or actual number of CPUs required. PBS runs the job when the specified resources are available, subject to constraints on maximum resource limits. There is more information on PBS resource limits.

Resources List Lets Users Choose

The different compute nodes have PBS resource names to identify the number of CPUs, memory, and CPU speed. It is important to specify the resources actually needed, which allows for better resource and job scheduling. The list below shows each compute node and resources available on that node:

```
davinci-01 sgi irix61B cpu1 mem128 mhz75
davinci-02 sgi irix61B cpu1 mem128 mhz75
davinci-03 sgi irix61B cpu1 mem256 mhz75
davinci-04 sgi irix61B cpu1 mem256 mhz75
davinci-05 sgi irix61B cpu8 mem2gb mhz90 mp
davinci-06 sgi irix61B cpu8 mem2gb mhz90 mp
davinci-07 sgi irix61B cpu8 mem4gb mhz90 mp
davinci-08 sgi irix61B cpu8 mem4gb mhz90 mp
```

These resources are defined as follows:

```
cpu1    = single CPU
cpu8    = eight CPUs
irix61B = SGI IRIX 6.1 BETA operating system
mem128  = 128 MB memory
mem256  = 256 MB memory
mem2gb  = 2 GB memory
mem4gb  = 4 GB memory
mhz75   = 75 mhz node
mhz9    = 90 mhz node
mp      = multiprocessor node
sgi     = Silicon Graphic PowerChallange node
```

The resource names allow users to select the nodes that best meet their requirements. For accounting purposes, users' node-hour allocations will be "charged" CPU time based on the wall-clock time for the job multiplied by the number of CPUs (not nodes) used. Thus, an 8-CPU node is eight times more "expensive" to use than a 1-CPU node. For detailed accounting information, see "System Accounting on the Cluster."

While it isn't necessary to specify resource requirements, it may well be an advantage to do so. For example, if the type of node isn't specified, PBS will assign the first available node(s) -- including the "CPU-expensive" 8-CPU nodes, which may not be what the user intended.

At times it is necessary to actually log into a node -- for example, to run a debugger. NAS prohibits users from logging directly into compute nodes without first acquiring the node(s) through PBS. However, users can access these nodes through the "interactive batch" function of PBS. To schedule an interactive session, use the qsub command with the "-I" option. PBS then provides a login shell and a list of nodes reserved for that job. From that shell, the user can do debugging (or whatever is necessary) and log directly into the nodes allocated for that job.

Request Help Via Email

For more information and assistance from the NAS parallel consulting staff, send email to **support@nas.nasa.gov**. A consultant will then contact you to answer questions and help solve problems.

For special requests, such as dedicated time on the entire cluster, send email to **davadmin@nas.nasa.gov**. Include the time and resources needed, along with a justification for the time requested.

To request an account on the NAS parallel cluster, contact Pat Elson, NAS user interface manager. Send email to **pelson@nas.nasa.gov** or call (415) 604-4463.

James P. Jones is the lead system administrator on the workstation cluster in the NAS parallel systems group. He is responsible for user accounting and utilization metrics on the NAS parallel systems, and is currently developing a C-language job scheduler for PBS on the cluster and on the SP2. He has worked at NAS for four years, and holds a Bachelor of Science degree from Santa Clara University, Santa Clara, CA.

Next Article | Contents | Main Menu | NAS Home

Machine	CPU	Memory	Network	Use
davinci	4	512 MB	Ether/FDDI/HiPPI/ATM	front end
davinci-01	1	128 MB	Ether/FDDI/HiPPI/ATM	compute node
davinci-02	1	128 MB	Ether/FDDI/HiPPI/ATM	compute node
davinci-03	1	256 MB	Ether/FDDI/HiPPI/ATM	compute node
davinci-04	1	256 MB	Ether/FDDI/HiPPI/ATM	compute node
davinci-05	8	2 GB	Ether/FDDI/HiPPI/ATM	compute node
davinci-06	8	2 GB	Ether/FDDI/HiPPI/ATM	compute node
davinci-07	8	4 GB	Ether/FDDI/HiPPI/ATM	compute node
davinci-08	8	4 GB	Ether/FDDI/HiPPI/ATM	compute node

Improving Job Performance and Turnaround on the CRAY C90s

Nicholas Cardo, Daniel DePauk, Robert Garza, and Alan Powers (NAS high speed processor group) identify two areas that can improve job performance and job turnaround time. The group continuously monitors the NAS and ACSF [Aeronautics Consolidated Supercomputing Facility] CRAY C90s, looking for ways to improve performance for the systems and for users' jobs.

Customizing Heap Space Size

When a user's job is running, the CRAY C90s maintain several areas of process memory, which assist processes in running. One of these areas, called "process heap space," is used for assigning more memory to a process quickly when needed. If the heap space is not large enough to satisfy the processes's allocation request, and there is not enough free system memory adjacent to the process, then that process is swapped out of memory.

When swapping occurs, both system overhead and job wall-clock time increase. On both the NAS CRAY C90 (vonneumann) and the ACSF C90 (eagle), the initial default heap size is 8,659 words. Increments of 4,096 words are added automatically, as needed.

In C language, process memory allocations occur when the routine "malloc" is called. In Fortran, these allocations occur when dynamic common memory is initiated by the first call in the defining subroutine. Due to the job memory demands on both C90s, if heap size needs to increase, the process will probably be swapped out. If a process needs to swap several times for memory allocation, wall-clock time increases significantly. This can be avoided by changing the heap management defaults when loading a program.

For example, the HSP group recently identified a user's process that was having problems with heap space. By changing the heap values, the process was shortened by 1,800 seconds of wall-clock time, and overall system overhead was decreased. This was accomplished by adding the following string to the Fortran compiling statement "cf77" when loading the program:

```
-Wl"-M fk7.map,s -H 236000000+100000"
```

The "-M" option produces a load map containing the sizes of the executable code, stack space, and heap space. The load map is useful for the "-H" option, which sets the initial heap size parameters and the increment size for heap space to expansion. The load map can also be used with the "-S" option, which sets the process stack space in a similar method.

Using dmget Adds Efficiency

Another way to decrease job wall-clock time is to use the dmget command for file retrieval. When a job accesses a file that is archived to tape, that job must wait while the files are restored to disk. Accessing several migrated files can add significant wall-clock time to the job.

The HSP group conducted a file retrieval study to show the various wall-clock times on both C90s when accessing migrated files. A test scenario was set up with six files, totalling 89 megabytes of data, that were archived to two tapes. Files 1-3 were on one tape, files 4-6 on another. Using four methods, the files were then accessed and retrieved via jobs that used dmget commands. In the first test, the files were retrieved one at a time alternately between the two tapes, forcing a tape mount for each file. With each file, the job waited for the tape to load, located the file on tape, and loaded the file onto disk. The total time to recall all the data is the cumulative time for each tape mount. This scenario yielded the worst performance results. (Opening files individually within a program without recalling the data in an optimum manner gives the same result.)

A slight performance improvement was achieved in the second test, when the files were sequentially recalled first from one tape, then the other. This test used a feature within the Cray Data Migration Facility (DMF), which keeps the tape mounted on a tape drive for a short period, in anticipation of sequential accesses to a single tape. This eliminates the time to mount the tape for subsequent accesses.

In the third test, all three files from one tape were accessed at one time, followed by all three files from the other tape. DMF can recall multiple files from the same tape in a single pass. This substantially improves the overall time in recalling files to disk.

The last test, in which all six files were accessed in one request, is the most efficient way to retrieve files. DMF can mount multiple tapes simultaneously, saving even more time.

The table shows the timing results from the four test scenarios.

An improvement of 8.5 minutes was gained by a simple change in the file retrieval method. To use this method, specify all the filenames required in one dmget command, for example:

```
dmget file.1 file.2 file.3 ... file.6
```

or, simply:

```
dmget file.[1-6]
```

The recommendations given here are simple and easy to perform and the end result may yield significant improvement in job performance and turnaround time, along with a decrease in system overhead -- a plus for all users.

Next Article Contents Main Menu NAS Home

	Test 1	Test 2	Test 3	Test 4
Commands	dmget file.1	dmget file.1	dmget file.[1-3]	dmget file.[1-6]
	dmget file.4	dmget file.2	dmget file.[4-6]	
	dmget file.2	dmget file.3		
	dmget file.5	dmget file.4		
	dmget file.3	dmget file.5		
	dmget file.6	dmget file.6		
Total Time	10.0 minutes	8.0 minutes	3.0 minutes	1.5 minutes

NASA's HPCC Program in 'Excellent Shape' Says Independent Review Team

by Larry Hofman

An independent review of NASA's High Performance Computing and Communications (HPCC) Program, held at Ames Research Center on June 6-7, reported that the overall program is in "excellent shape."

The findings of the Independent Annual Review (IAR) panel were presented to NASA Administrator Daniel Goldin and Associate Deputy Administrator Jack Dailey on July 18. The report stated, "The HPCC Team competence is very high, well managed, and project working relationships appear to be excellent." It added that the HPCC Program "showed substantial progress in all areas" since last year and is on track with objectives and plans.

NASA's major programs are reviewed annually by teams consisting of officials from NASA and other government agencies. The IAR panel, which included representatives from NASA headquarters, Johnson Space Center, Marshall Space Flight Center, Jet Propulsion Laboratory, Wright Laboratories, and the Department of Energy, assessed HPCC's progress and performance against planned milestones and verified its ability to meet future commitments.

The review consisted of presentations by HPCC project managers, extensive discussions based on these presentations, and telephone interviews by panel members to selected customers, including McDonnell Douglas Aerospace, Pacific Gas & Electric, and the University of Washington. These interviews rated HPCC's directions and products very positively, according to the IAR panel's executive summary.

Responding to Customer Feedback

During the review William J. Feiereisen, Computational Aerosciences (CAS) Project Manager, noted several changes made to the CAS Project during the past year in response to evolving customer requirements. A near-term goal of cost-effective, high-performance computing (in the form of clustered workstations) was added to the original, longer-term goal of extreme high-performance (teraflops) computing. The set of Grand Challenge problems, which CAS uses to represent design applications that drive the need for commercial aerospace computing, were shaped to be more relevant to industry needs.

Feiereisen stated that U.S. industry and academia now have more direct involvement in CAS through a newly formed Review and Planning Team, whose purpose is to obtain guidance and feedback from industry and academia. Another example of direct involvement is through the use of cooperative research projects. He noted that 50 percent of CAS technology development and demonstration is now performed by these two sectors (in cooperation with NASA), which helps assure technological relevance and facilitates technology transfer.

Significant Accomplishments for 1995

High on the list of CAS accomplishments are successes with the Cooperative Research Agreement (CRA) that provides IBM SP2 systems at NAS and Langley Research Center. The 160-node system at NAS achieved production reliability ahead of schedule and is operating with more than 65 percent CPU utilization. In April, 144 nodes of the SP2 were used by Charbel Farhat (University of Colorado, Boulder) to achieve sustained performance greater than 29 gigaflops for a large, out-of-core problem. (See NAS News May-June '95 for more information about Farhat's calculation.)

Clustered workstation testbeds, each with different configurations, are now in operation at Ames, Langley, and Lewis Research Centers in support of the new cost-effective, high-performance computing initiatives. The most powerful of these, the cluster of 16 Silicon Graphics Inc. Power/Challenge L workstations at Ames, is available for 24-hour dedicated use following a difficult six-month integration effort. The NAS-developed Portable Batch System (PBS) was installed on the SP2 and the Power/Challenge L cluster, with throughput and CPU utilization increases of more than 100 percent.

Associated with the clustered testbeds research, an important CRA for Affordable High Performance Computing was signed in May with a consortium of companies and universities led by United Technologies Corp.'s Pratt & Whitney Group. Other new contracts include seven NASA Research Announcements, which were signed with various university and industry participants.

Brief History of HPCC Program

HPCC, a federal program with an FY95 budget of nearly $1.04 billion, involves nine government agencies, of which NASA is a major participant. The program, initiated in 1991, fosters development of advanced, high-risk, high-payoff computer and telecommunications technologies to benefit the United States' competitive position. NASA contributes research in all four HPCC project areas: Computational Aerosciences (CAS), Earth and Space Science (ESS), Information Infrastructure Technology and Applications (IITA), and Remote Exploration and Experimentation (REE). Ames provides overall management of NASA's CAS and IITA projects, which account for much of the $76 million NASA HPCC budget. The NAS Systems Division has a lead role in several elements of the CAS project, including computing testbeds, system software and tools development, and the National Research and Education Network (NREN) Program.

Larry Hofman is a member of the CAS project management team.

Next Article | Contents | Main Menu | NAS Home

Next Article | Contents | Main Menu | NAS Home

SP2 and `davinci' Used for Ames MAPPS Project Test

by Elisabeth Wechsler

The High Performance Computing and Communications Program's (HPCCP) 160-node SP2 ("babbage") and 8-node, multiple CPU Silicon Graphics Inc. PowerChallenge L workstation abstracts ("davinci") at the NAS Facility were connected by a dedicated FDDI network to the 40x80 Ames wind tunnel to perform noise tests August 14-31 on a McDonnell Douglas DC-10 aircraft model.

The test was conducted using MAPPS [Microphone Array Phased Processing System]. "MAPPS takes data from 40 microphones in a 4-ft. sq. pattern, processes 200 frequencies, and returns results within five minutes, allowing us to determine noise locations and intensities more accurately during wind tunnel tests," explained Mike Watts, aeronautics engineer with Code IC and MAPPS team lead.

The SP2 was selected for "fast turnaround," according to Watts, and also for its potential in prototyping future measurement systems. davinci, which served as test backup for the SP2, performed calibrations for the pre-test work.

The SP2 and davinci comprise two of the seven elements of MAPPS' hardware and software system. The MAPPS project, part of the NASA Advanced Subsonic Transport noise-reduction program, represents a joint venture between HPCCP, DARWIN, and Ames' Computational Sciences Division and AA divisions. For more information, send email to **watts@ptolemy.arc.nasa.gov**.

New Technical Reports

Summaries of some recently published technical papers written by NAS staff members are shown below:

JSD: Parallel Job Accounting on the IBM SP2

William Saphir, James Patton Jones

The IBM SP2, one of the most promising parallel computers for scientific supercomputing, is fast and usually reliable. One of its biggest problems is a lack of robust and comprehensive system software. IBM's sortware doesn't provide accounting for parallel jobs, other than what is provided by AIX for individual process components.

Without parallel job accounting, it is not possible to monitor system use, measure the effectiveness of system administration strategies, or identify system bottlenecks.

To address this problem, we have written *jsd*, a daemon that collects accounting data for parallel jobs. *jsd* also notifies system administrators in certain cases of system failures.

Technical Report NAS-95-016

CFD Data Sets on the WWW for Education and Testing

Al Globus

The NAS Systems Division at NASA Ames Research Center has begun the development of a Computational Fluid Dynamics (CFD) dataset archive on the World Wide Web.

The datasets are integrated with related information such as research papers, metadata, and visualizations. The author identifies and discusses four classes of users: students, visualization developers, CFD practitioners, and management. Bandwidth and security issues are briefly reviewed, and the status of the archive as of May 1995 is examined. Routine network distribution of datasets is likely to have profound implications for the conduct of science. The exact nature of these changes is subject to speculation, but the ability for anyone to examine the data, in addition to the investigator's analysis, may well play an important role in the future.

Technical Report NAS-95-015

NAS Focuses on Virtual Reality and Web Tools at SIGGRAPH

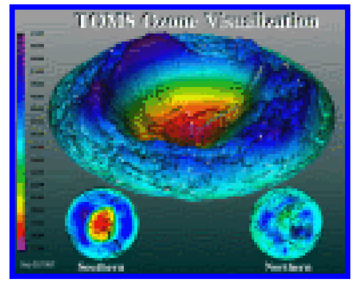

by Jean Clucas

"Interactive" was the word at SIGGRAPH '95, held August 6-11 in Los Angeles. New exhibits this year were Interactive Communities, which demonstrated collaborative projects using networks, interactive media, and graphics; and Interactive Entertainment, which featured games and virtual reality (VR).

NAS participation in these hot topics included the course "Developing Advanced Virtual Reality Applications," chaired by Steve Bryson, who developed the NAS Virtual Windtunnel. Bryson was also a panel member for "Performing Work in Virtual Environments." Both sessions explored how to make virtual environments useful.

Quick Tips for `Effective VR'

Bryson said that, for effective VR, he had learned to use opportunistic techniques such as intrinsic metaphor -- for example, walking through a building within an architectural VR model.

Embedding the user interface in the VR environment is also important. Frequently, the interface should be part of the metaphor, such as pointing to and grabbing objects. Bryson suggests using a few gestures only and including buttons when appropriate. He noted that buttons don't need to follow a real-world metaphor -- for example, buttons can "float" in space, even though that isn't realistic.

FASTexpeditions in `Shoot Out'

The panel session "3D Graphics Through the Internet -- A 'Shoot Out'" pitted FASTexpeditions, designed by Val Watson, of the NAS information management group, against VRML browsers, Apple Computer Inc.'s QuickTime VR, and web server-controlled visualization.

Watson believes that FASTexpeditions, tools for automatically running datasets from the web in FAST, are preferable for scientific data because they transfer real data that can be visualized with scripted tours, then analyzed further by users, and viewed collaboratively at multiple sites.

See also FASTexpeditions into EPA Pollution Data -- Ozone in the Northeastern U.S.

Sept - Oct 1995, Vol. 2, No. 12

Executive Editor: Marisa Chancellor

Editor: Jill Dunbar

Senior Writer: Elisabeth Wechsler

Contributing Writers: Nicholas Cardo, Jean Clucas, Daniel DePauk, Robert Garza, Larry Hofman, James P. Jones

Image Enhancements: Chris Gong, John Hardman

Other Contributors: John Adamczyk, Mark Barnett, Al Bayucan, Chris Beaumont, Walt Brooks, Steve Bryson, Chris Buchanan, Mark Celestina, Dennis DeRyke, Pat Elson, Dave Faust, William Feiereisen, Al Globus, Karen Gundy-Burlet, Lisa Harris-Calhoun, Vee Hirsch, Mary Hultquist, Chris Kleiber, Amy Lacer, Tom Lasinski, Creon Levit, Kevin McCabe, Richard Mulac, George Myers, Alan Poston, Alan Powers, Marcia Redmond, Dani Thompson, Mike Watts.

Editorial Board: Marisa Chancellor, Jill Dunbar, Chris Gong, Eric Hibbard, Kristina Miceli, Chuck Niggley, Serge Polevitzky, Pamela Walatka, Elisabeth Wechsler

NAS NEWS

Volume 2, Number 12 September – October 1995

Large-scale Model for Analyzing Aircraft Engine Design Challenges Supercomputers at NAS Facility

by Elisabeth Wechsler

The supercomputers housed at NAS—eagle (CRAY C90, 8 CPUs), vonneumann (CRAY C90, 16 CPUs), and babbage (IBM SP2, 160 nodes)—are being used this summer to run a quasi-production code that attempts to analyze the new components of an aircraft engine.

The mathematical model, referred to as the average-passage flow model, describes the three-dimensional flow in multistage turbomachines, and was developed by John Adamczyk, senior aerospace scientist, and his team at NASA Lewis Research Center.

This work has had "a significant impact on the PW4000 High Pressure Compressor development effort" at Pratt & Whitney, a United Technologies Corp. subsidiary, according to Mark Barnett, senior research engineer at United Technologies Research Center. The model allows engineers "to obtain much more accurate predictions of the performance and characteristics of their multistage turbomachinery designs," he added.

So far, the NAS Facility's SP2 has been used to run a simulation for 17 (out of a possible 23) blade rows of a Pratt & Whitney high pressure compressor (HPC). Adamczyk's goal is to run the full 23 blade rows in parallel "by the end of the summer."

Reduces Design Time, Cost

The use of Adamczyk's model within Pratt & Whitney's HPC analysis code has contributed to a 2 percent performance improvement over the baseline design with increased stall margins, and helped to reduce the design time from two years to one.

The potential cost savings could be "in the tens of millions of dollars" through reductions in the number of test rigs that would need to be built and tested, Barnett commented.

"Engine manufacturers are waiting in line to use this code," said Jim Cross, Acting Chief of the

Continued on page 2

This Issue:

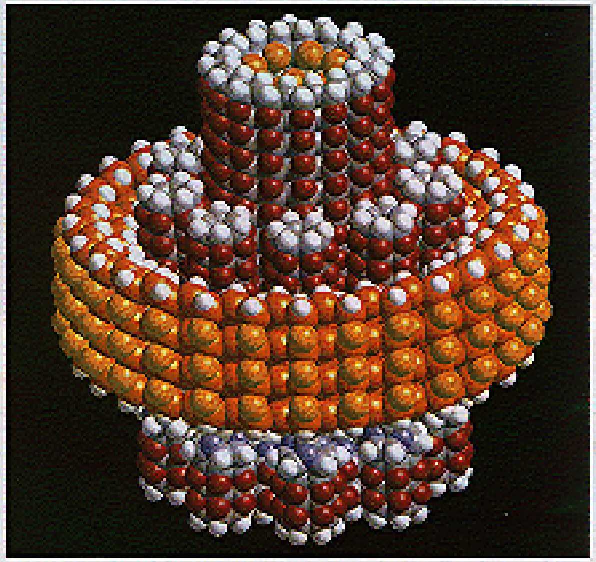

Molecular model of a hypothetical planetary gear, used in mechanical engineering, in which the spheres represent individual atoms. This simulation suggests that moderately complex mechanical parts can be built from precisely positioned atoms. The image was designed and simulated by K. Eric Drexler (Foresight Institute, Palo Alto, CA) and Ralph Merkle (Xerox Palo Alto Research Center). It was rendered using Molecular Assembly Sequence Software (MASS), developed by Carol B. Shaw. For more information, see the World Wide Web at http://www.portal.com/~carol/mass.html

Nanotechnology Creates New Opportunities for NAS, Ames, Industry

by Elisabeth Wechsler

Molecular nanotechnology—the building of products by placing atoms in precise locations using controlled trajectories—is "a promising future technology," said David Bailey, NAS senior scientist, in his introduction to a New Technology seminar on this subject at the NAS Facility, on July 26.

Scientists Al Globus and Creon Levit, of the NAS information management group, told the audience of some 80 NASA Ames researchers that the use of nanotechnology could achieve "orders-of-magnitude improvements in aerospace vehicles and computers." The speakers illustrated presentations of what they envision nanotechnology can offer the aerospace industry—and the role NAS could play.

Bailey, Globus, and Levit acknowledged the "substantial risk and long lead time (5-20 years)" for this field, but they nevertheless believe nanotechnology "appears feasible" and recommended that NAS "aggressively seek a dominant role in computational nanotechnology."

Computational Nanotechnology

One area particularly well suited to NAS is the "model" area of computational nanotechnology, in which molecular editing, simulation, modeling, testing, and validation are used to design "molecular products with millions of atoms and thousands of assembly steps," Levit said.

According to Bailey and Globus, potential applications of computational nanotechnology include:

- ultrastrong components for jet engines
- ultralight materials for space vehicles
- computer processors with dramatically higher performance and lower energy requirements
- computer memory devices with dramatically higher storage densities

Levit explained that physical experiments for nanotechnology are often impossible and that simulation is necessary to establish design.

Continued on page 2